Understanding IT Service Management with Scianta Analytics Extreme Vigilance

Introduction

Business support services like payroll and expenses are already difficult to monitor and manage without service level agreements, but when end users start filing tickets the pressure goes up. The hard-to-manage realities of service management and incident lifecycle is extremely well-suited to Scianta Analytics' Cognitive Computing approach. Seasonality, change freezes, and continually shifting regulatory environments make it tough to develop sensible monitoring and alerting systems that are responsive to user complaints. Scianta's Extreme Vigilance products are ideally positioned to assist the operational support team with solving common business-facing problems.

Elastically scaling service architectures built on virtual machines or containers are changing the landscape of operational services provision, moving the point of monitoring attention away from servers and towards services. Such a transition is even more pronounced for organizations that are building on SaaS, PaaS, or IaaS in the cloud and may not be able to control monitoring directly. This change in monitoring focus necessitates revisiting the prioritization of infrastructure oriented approaches, and increasing the use of behavioral analytics and synthetic monitoring.

Problem Statements

Alerting as a Trailing Indicator

🚨 🚨 🚨 🚨 🚨 🚨 🚨 🚨 🚨

No one maintaining a service for their customers wants to see problem reports! Perhaps there are known key performance and security indicators to watch, or perhaps they can be discovered. Even so, monitoring dozens to hundreds of indicators for a change that might be important is ridiculously expensive. Avoiding problems costs a lot less than fixing them, which is great if you know about problems before they have occurred. Deviation-from-norm alerts are a good starting point, but it takes a much deeper contextual awareness of transactional patterns and seasonality to avoid predictable alert storms.

Overwhelmed Analysts

😔 😕 🤕 😭 😔 😳 😐

One of the worst feelings is to discover a problem with a long tail of old unresolved incidents that no one had recognized. Users file incidents with the symptoms that they observe, but the imaginative work required to grasp how these symptoms might occur is blocked on an overloaded individual or lost in the noise of more easily solved problems. Analysis of the incidents that are filed can uncover useful patterns that may elevate groups of these long tail tickets automatically.

Where's the Data?

👓 👓 👓 👓 👓 👀

Data acquisition is also a significant challenge in many environments. Service or software providers may not have prepared for data export, or may only offer reports as a limited afterthought. Exotic APIs, rigid third-party support contracts, and non-negotiable regulatory commitments mean that data quality challenges and instabilities are built in. Alternate means can be used to collect data in some organizations, such as passive network capture or access to a database backend managed by another team or vendor; but in other organizations the data scientist can only access extremely limited data sets.

Sound and Fury

🎶🎵 🔥🔥 🔥🔥 🎵🎶

Monitoring is hardly new, and in some cases multiple regimes have come and gone; often they're removed because of the noise generated. While alert noise is arguably better than nothing, a monitoring system that disrupts and exhausts analysts is not helpful. Analysts expect monitoring systems to weigh risk and context appropriately, respond smartly to change, and learn from correction. Above all, the reasons for an alert generation must be clearly discoverable so that logical mistakes can be corrected.

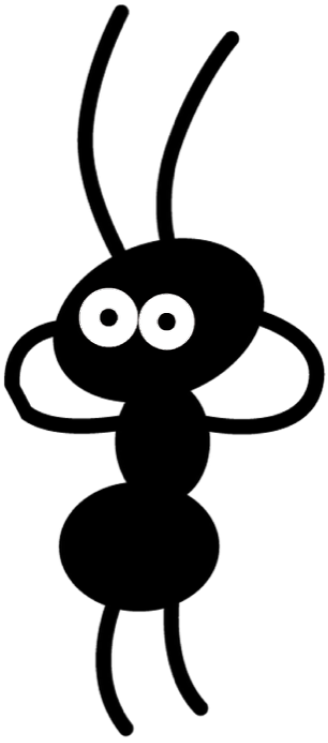

Suggestions

Alerting as a Trailing Indicator

Extreme Vigilance can model the transactional behavior of an Actor. While metrics and counts may be used to indicate technical problems, alterations in the way that people and systems interact are indications at a business level. Transactional analysis plus anomaly detection opens the door to a better understanding of operations across the board. This approach is particularly valuable for service monitoring, in which a service like OCR receipt scanning is the actor and the containers on which it runs are just attributes of an action. By keeping focus on the actor's activities instead of the health of servers, the amount of noise is reduced and the accuracy of alerting is increased.

Extreme Vigilance also improves on the basic anomaly detection capabilities of more traditional monitoring systems. The analyst and data scientist work together to define a Data Dictionary which describes the events of interest in terms of actors, assets, and actions. These events are reviewed in a Cognitive Model which automatically determines the band of normalcy for each combination of actors and assets by specific time frames. As new values arrive in the data stream, Extreme Vigilance qualitatively measures the fit with observed data, emitting signals when measurements are approaching or breaching calculated thresholds.

These emitted signals can be used to trigger incident alerts of course, but analysts are also able to define crisp rule sets in the Cognitive Rules Engine to describe known compliance issues and take contextual facts into account. For instance, a rule could be written to state that the maximum dwell time for a reported incident must not rise above 12 working hours. Additionally, rules could be written to express growing concern when the overall dwell time is trending upward, or trending upward rapidly, or trending upward very rapidly. Extreme Vigilance calculates reasonable meanings for "very" and "rapidly", while respecting the crisp limit of 12 hours.

Another excellent way to alert ahead of issues is to use Actor and Peer Analyses. These techniques review recent behavior of resources in order to determine how well the behavior matches with past behavior or the behavior of similar resources. If recent measured values are anomalous, signals are emitted for analysis. Comparing an Actor to itself or to its cohort uncovers subtle variations that may not be visible in threshold anomaly monitoring due to a gradual slope. For instance, peer analysis of filed incidents relating to an expense reporting system may be used to uncover a currency conversion problem, while transactional analysis of the expense reporting tool could be used to find that problem's root cause.

Overwhelmed Analysts

Peer and Actor analysis of filed Incident tickets in the service desk is an extremely interesting technique. As a time-series index, Splunk and Scianta inherently work together to detect anomalous spikes in ticket filing when sorted by user groups, time frames, or affected services. Additionally, whether using basic Splunk pattern matching or an advanced Natural Language Processor like Insight Engines, administrators can review

free form fields in tickets for keyword matches that may indicate known symptoms.

Where's the Data?

👓 👓 👓 👓 👓 👀

Missing data is an insurmountable problem for many data analysis systems, leading to accuracy challenging techniques like interpolation and synthetics. Scianta's use of the Splunk platform makes it possible to glean high amounts of value from incomplete, indirect, and disparate data streams. Splunk's rich add-on ecosystem enables access to raw packet capture, infrastructure logs and metrics, database tables, and APIs. While direct access is certainly preferred, Splunk makes it possible to monitor a system indirectly via its impact on infrastructure. In turn, Scianta's behavior analytics can then operate on these indirect signals.

Sound and Fury

🎶🎶 🔥🔥 🔥🔥 🎵🎶

Scianta Extreme Vigilance puts significant effort into producing quality alerts. Signals are generated from a wide variety of rule matches, anomaly detections, and analysis results, but these signals are all weighted. Signal intensity weights are scaled by the severity of the breach, trust level of the model, and criticality of the resources. Analysts can then dial the system's propensity to alert up and down based on their resources and trust in the system's accuracy.

Conclusion

Scianta's Cognitive Computing approach provides transactional analysis and anomaly detection as augmentation to human analysis, purely within the existing Splunk environment. To learn more, see https://www.scianta.com